Sentiment of the FOMC: Unscripted

By San Cannon

The Federal Open Market Committee (FOMC) meets eight times each year to set monetary policy. During these meetings, a changing cast of participants engages in presentations and discussions, drawing on the perspectives of research staff and community and business leaders as they formulate their views on economic conditions and determine the stance of monetary policy.

Determining what the FOMC finds relevant to policy discussions and how these discussions might have changed over time can be challenging. Although the Committee releases carefully constructed statements and meeting minutes to the public, some marketwatchers have argued these pieces have only rendered proceedings more mysterious or opaque. The full transcripts offer a more complete picture of Committee meetings; however, these transcripts are only released to the public after five years. Furthermore, the transcripts can be somewhat difficult to parse: the texts contain a wealth of disparate information ranging from casual anecdotes to research findings to staff economic forecasts.

Nevertheless, meeting transcripts offer readers the unique opportunity to examine the original expressions of individual meeting participants prior to being distilled and summarized into the statement and minutes. Applying text-mining techniques to FOMC transcripts can help quantify this information to provide a rich analytical resource reflecting real-time economic and financial analysis. The words

---

San Cannon is an assistant vice president and economist at the Federal Reserve Bank of Kansas City. This article is on the bank's website at ***www.KansasCityFed.org***

participants choose for particular topics allow text analysts to measure the tone of the overall discussion in a way not possible in statements or minutes. In addition, researchers can measure the tone of individual speakers. Unlike the minutes, which attribute general summary discussions to unidentified "Committee members," the transcripts identify speakers along with their contributions. This identification invites comparisons between individual speakers or classes of individuals such as Board members and Bank presidents.

In this article, I study the tone and diction, or word choice, of the meeting participants to better understand how the discussions are formed, how they related to the performance of the economy, and how they may have changed with movements toward greater transparency. Using some fairly simple language-processing tools, I measure the tone of FOMC deliberations, explore differences across speakers, and examine how the tone of the discussions relates to a measure of economic activity. I find first that the composition and tone of the discussions have changed over time. More specifically, the length of comments, the uniqueness of word choice, and the measure of the tone display distinct patterns from the late 1970s through 2009. Second, I find measurable differences in the diction and tone of different classes of speakers who participate in the discussions. The contributions of Board members, for example, have a different composition and tone than that of Reserve Bank presidents or Federal Reserve System staff. Finally, I find measures of the relationship between the tone of the discussions and economic activity also show differences across time and speaker.

Section I provides background information on the transcripts and the text-mining tools used to extract information. Section II calculates the tone measure for each discussion and explores how the role of individual speakers has changed over time. Section III examines the relationship between the tone measure and real economic activity and assesses what effect a move toward greater transparency in the Committee might have had.

I. Extracting Text from the Transcripts

Committee discussions generate an extensive amount of text. Although the Federal Reserve Act only mandates four FOMC meetings per year, the Committee met as often as monthly up until the early 1980s and has met eight times each year since. Conference calls may also occur between scheduled meetings. In addition, the number of meeting attendees contributing to the deliberations can add significantly to the text. The Committee comprises all sitting members of the Federal Reserve Board of Governors—usually seven but at times as few as four—as well as five Reserve Bank Presidents who serve on the Committee on a rotating basis. Reserve Bank presidents who are not voting members of the Committee attend and participate in all meetings as do staff members from Reserve Banks and the Board. The meetings are closed to the public, but the Committee releases an official statement at the close of the meeting to convey its monetary policy decision. Minutes from the meeting are available several weeks later, and the Committee releases full transcripts of the discussion with a five-year lag.[1]

Not all of these communication pieces may be suitable for text analysis. The official statements, for example, are perhaps too carefully crafted, as the media and market participants vigilantly parse them. Indeed, *The Wall Street Journal* dedicates a column to outlining changes in the wording of FOMC statements from meeting to meeting. The transcripts, on the other hand, are ideal for text analysis, as they capture each part of the meeting from roll call to parliamentary procedures for policy votes. The transcripts include the entire discussion, indicating who was speaking and what was said with little editing except for the potential removal of "a very small amount of information received on a confidential basis from, or about, foreign officials, businesses, and persons that are identified or identifiable" (Board of Governors 2014). They show how Reserve Bank Presidents provide important regional context and information, how Governors voice opinions or ask questions, and how Board staff present information on economic output and other relevant topics. Such detail makes the text of the transcripts an excellent source of information to be mined.

Text mining

Text mining creates structured data out of unstructured data, allowing a quantitative analysis of qualitative information. Traditional methods of assessing relationships and patterns in data deal exclusively with structured data—numeric information generally well-formatted in tables or databases. However, much of the data created or captured today is far less structured or in many cases unstructured, such as the text of tweets, blog posts, emails, or documents. Analyzing such inputs first requires transforming them from the raw data format to a format that can effectively use methods identical or analogous to those used to analyze structured numeric data.

Researchers can mine text using different methods, each suitable for answering a different set of questions. More specifically, new research in this field applies a variety of methods to FOMC documents.[2] I focus here on the specific words FOMC participants choose during their discussions. Note that studying the words used is different than examining the topics discussed. The former is more closely aligned with expression, the latter with content.

As expression and word usage relate more to how ideas are conveyed than to the ideas themselves, they are a more appropriate way to address sentiment in a document like the FOMC transcripts. To assess how someone feels, examining their actual choice of words may be more instructive than attempting to attach a sentiment to a particular topic. Much of the current work on sentiment analysis focuses on consumer opinions expressed in tweets, online reviews, and other social media outlets. I apply similar techniques here with some changes to acknowledge the important differences between social media posts and monetary policy discussions.

Processing the transcripts

Some written records of the Committee's meetings are available from the Federal Reserve Board from as early as 1936. I start the sample with 1977, as this is the first year for which records are identified as transcripts. First, I extract the text from the digital file, parse it into words based on spaces and punctuation, and remove the preliminary Committee procedures (for example, roll call). I then group the text pieces into individual comments by speaker. For each named speaker, I collect the text of that person's comment until the next speaker is identified. For some entries, this text is as short or simple as "yes" or "thank you"; for others, a speaker giving a presentation or answering a question at length can have a single comment that runs for pages. I apply the extraction method to 362 complete transcripts and five partial transcripts over 33 years, yielding 114,912 individual comments.[3] I

Chart 1
Distribution of Transcripts over Time

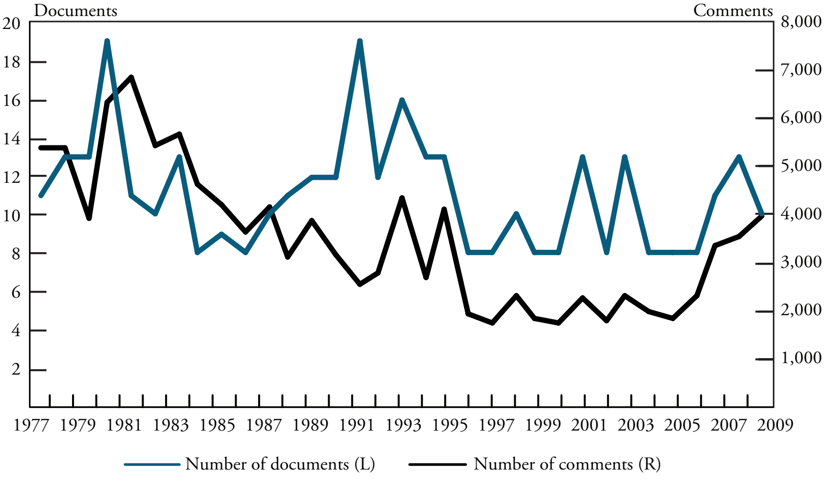

Sources: FOMC and author's calculations.

remove numbers and punctuation from each comment, and convert all words to lower case to facilitate matching with words which may have different capitalization.

The number of transcripts and the number of comments extracted from the transcripts has varied drastically over the sample period. Chart 1 shows the distribution of the number of documents and the extracted comments over time. The FOMC convened as many as 19 times in 1980 (10 meetings and nine conference calls) and 1991 (eight meetings and 11 conference calls) and held the current standard number of meetings (eight) in 10 of the 33 years.

The next step in processing is to eliminate what are known as "stop words": common words such as articles ("the"), conjunctions ("and"), and helping verbs ("would," "are") unlikely to reveal any interesting information when examined thoroughly.

Even without these common words, the meeting participants had plenty to say. The parsed transcripts contain 4,746,165 words, after excluding 2,731,724 instances of 100 stop words. Figure 1 displays the most common 100 words from the 29,802 different words in the transcripts. Had I not removed the stop words, the top five most common words would be "the," "that," "and," "have," and "are," which don't

Figure 1
Distribution of Top 100 Words

Note: Size of the word indicates the relative distribution of that word.
Sources: FOMC and author's calculations.

contribute much insight into the Committee's discussion. As Figure 1 shows, eliminating stop words does not eliminate neutral or uncommunicative words entirely. In addition to spoken words, the transcripts include some descriptors that the word count also captures. In 2,369 instances, the commentary was "unintelligible," and in 3,225 instances, "laughter" was documented.

Word frequency should not be strictly interpreted as indicating the importance or relevance of particular topics, as speakers may use synonyms and more detailed descriptors. Simple word counts are included here to give some insight into how Committee members most often express particular ideas and topics, which is part of the tone assessment discussion that follows. That said, Hansen, McMahon, and Prat apply a topic extraction method to the transcripts and show that for some topics, a single word or two clearly dominates all others in the discussion. In their analysis, the words "inflation" and "unemployment" are

Chart 2
Number of Comments per Speaker per Year

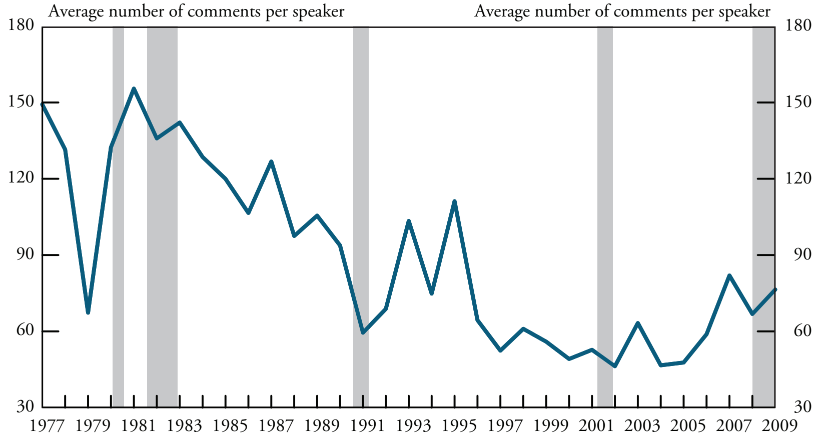

Note: Gray bars represent NBER-defined recessions.
Sources: National Bureau of Economic Research, FOMC, and author's calculations.

associated with the same topic or concept. Indeed, the two are inextricably linked for much of the FOMC's focus, as both are part of the Fed's "dual mandate" of maximum employment and price stability. But the word "inflation" overwhelmingly dominates all other words associated with what they label the "inflation" topic, including "measure," "core," and "percent."

While many pre-processing options are available in different text-mining applications, I choose to minimally pre-process the text. Removing stop words, for example, is helpful for exercises involving word counts or relative frequencies but may not be helpful for other analyses. In addition, I have chosen not to weight the words when calculating the sentiment measure. None of the commonly used weighting schemes is an obvious choice for this exercise, and though some evidence suggests weights can help decrease the noise in certain measures, it is not clear they would improve this analysis.

Simply counting the number of comments in the transcripts highlights changes in the nature of the FOMC's discussion over time. Chart 2 shows the average number of comments made by each speaker each year during the meetings or conference calls that occurred that year. The number of individual contributions per speaker has varied greatly over time with a distinct downward trend through about 2005 and a

Chart 3
Average Number of Words per Comment

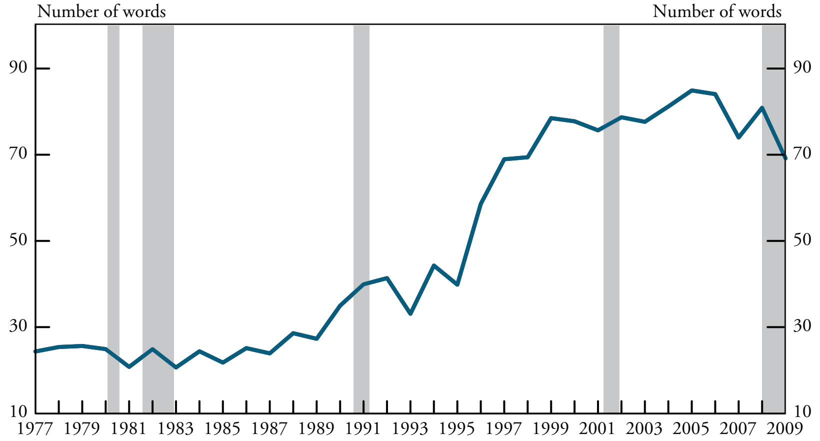

Note: Gray bars represent NBER-defined recessions.
Sources: National Bureau of Economic Research, FOMC, and author's calculations.

steady upward trend since then. There does not appear to be any clear cyclical pattern.

But fewer comments do not mean less discourse. Indeed, while the number of comments decreased, their average length increased. Chart 3 shows the number of words per comment appears to have increased significantly from 1993 to 2005. While the number of comments has decreased from its peak in 2005, it is still significantly higher than in the early years of the sample.

Although Committee members might use many words to discuss a certain concept, some words are core descriptors of monetary policy objectives and deserve individual attention. Chart 4 shows the distribution of particular words over time as a percentage of total words used. As other research has noted, the Committee uses the word "inflation" far more often than "unemployment" in its discussions. Members mention "growth" frequently as well, though they seem to rarely mention "employment." Hansen, McMahon, and Prat find "growth" and "employment" belong to two distinct topics separate from the "inflation" topic which encompasses the words "unemployment" and "inflation." The "growth" topic contains words such as "expansion" and "increase"; in contrast, the "employment" topic contains words such as "district"

Chart 4
Word Appearances per Comment

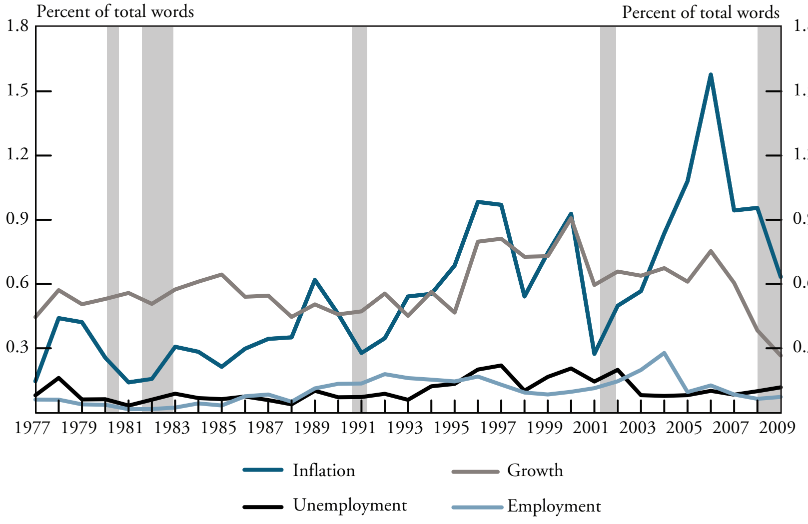

Note: Gray bars represent NBER-defined recessions.
Sources: National Bureau of Economic Research, FOMC, and author's calculations.

and "region," possibly indicating different focuses among sets of meeting participants. Members have used "inflation" much more in recent years: instances of the word peaked in 2006, when it made up more than one percent of all words in the Committee discussion.

II. Measuring Tone and Speaker Effects

Individual words can have a specific semantic orientation, meaning that they consistently convey a positive, negative, or neutral sentiment regardless of the topic with which they are affiliated. For example, "admirable" generally conveys a positive notion or idea. On the other hand, "lost" may more often have a negative connotation. Most words, though, have a neutral orientation: the word "word," for example, doesn't necessarily convey a positive or negative sentiment.

To measure the sentiment of FOMC discussions, I examine each comment, first evaluating the orientation of each word in the comment and then calculating a tone measure for the comment as a whole. Thus, the tone measure captures the net sentiment of the comment as either positive, negative, or neutral. For each month in which either a meeting or conference call takes place, I use the tone measures for each comment to calculate an overall tone metric as an indicator of sentiment for that month.

Creating the tone measure

The tone of a comment is determined by the semantic orientation (positive, negative, or neutral) of the words in that comment. How researchers determine a word's semantic orientation depends on the word list or dictionary they use to evaluate it. Many researchers have created their own dictionaries to evaluate tone. One common approach is to start with a set of seed adjectives that carry a clear semantic orientation and then augment that list by attributing the tone of a seed adjective to its known synonyms. Another method is to consider word classifications that researchers have created in other domains, such as psychology, and edit them to fit a particular use case.

One wordlist constructed using the boosting-by-synonyms approach is that of Hu and Liu, who have worked on opinion mining and sentiment analysis of online customer reviews, social media posts, and other Internet venues. Hu and Liu start with 30 seed adjectives and, using their synonyms, create lists containing 2,006 words with a positive orientation and 4,783 words with a negative orientation. While Hu and Liu apply these word lists to consumer good evaluations, the lists are general enough to be suitable for a broader use. The list of positive words, for example, is quite extensive, ranging from "cozy," "swanky," and "twinkly," which one may not expect to find in a monetary policy discussion, to "outperform," "judicious," and "insightful," which may be more likely candidates. The range of negative words is equally large, from less formal words such as "anarchy," "stupidity," and "zombie," to the more reserved "worthless," "sluggish," and "inflationary." Although they may be less formal words, "anarchy," "stupidity," and "zombie" are all found in the FOMC transcripts. Because this dictionary was compiled to evaluate customer ratings of consumer goods, I refer to it as the "consumer" dictionary throughout the text.

While its broad range of words makes this dictionary appealing, the set of general words may not be a good fit for the specialized content of the transcripts. An alternative approach is to consider a dictionary tailored more specifically to financial and regulatory discussions.

Loughran and McDonald start with word classifications used in psychology and construct a dictionary more suitable to classifying financial text. Using text found in Securities and Exchange Commission (SEC) filings from 1994 to 2008, they build a dictionary of 85,131 words classified in multiple sentiment categories: positive, negative, uncertain, litigious, and constraining, among others. Of the larger list, 2,355 words are identified as negative, 354 as positive, and 297 as conveying uncertainty. And unlike other dictionaries, sentiment assignment is not mutually exclusive: in multiple instances, words appear simultaneously on two lists, usually uncertain and negative. For example, "anomaly," "doubt," and "deviate" are deemed to convey both negative and uncertain sentiment.

Loughran and McDonald's motivation for compiling this specialized dictionary is similar to a concern faced in this article: specifically, the context for words in technical documents like financial filings may be different than for other text domains. Indeed, Loughran and McDonald find only half of the words on their negative list appear on an alternative general-use sentiment list. Even when they do appear on a general-use list, the sentiment of these words may differ significantly in a technical context. For example, "liability" in a financial filing is often used in an accounting sense rather than as a pejorative term. However, this technical context may not make the Loughran and McDonald dictionary a better fit for the FOMC transcripts. In its focus on financial filings, their "finance dictionary" excludes some more common, relevant words from the consumer dictionary (such as the previously highlighted words "outperform," "insightful," and "sluggish").

Using both the finance and consumer dictionaries to score the tone of a comment can show why dictionary choice is so important. Take, for example, the following quotation from former Chairman Alan Greenspan: "It's an interesting question: When does this long-term trend we are all forecasting begin to affect the M2 data?"

After processing to remove punctuation, numbers, and capitalization, this comment appears in our calculations as: "its an interesting question when does this longterm trend we are all forecasting begin to affect the data." Once the number 2 is removed, the standalone "M" isn't recognized as a word or noted as an appropriate abbreviation and so is also dropped from the processed text.

I can then calculate a tone score (positive, negative, or neutral) for this comment based on each dictionary's assessment of the words it contains. The calculation for the tone label is that used in Fuksa and Sornette, and Sadique and others, and is applied for each of the dictionaries:

Tone = (#positive words – #negative words)/(#positive + #negative),
Tone > 0 indicates positive tone,
Tone < 0 indicates negative tone, and
Tone = 0 indicates neutral tone.

The consumer dictionary gives Greenspan's comment a positive tone label, while the finance dictionary gives it a negative tone label. The difference in labels comes from the different classifications of individual words. "Question" is the only word in Greenspan's comment that appears in the finance dictionary; it has a negative label, so the comment gets a negative label. The only word that appears in the consumer dictionary is "interesting," a positive word, so the comment gets a positive label.

The tone of a word can be changed by the words around it. When strictly scoring words with a dictionary entry, it is easy to miscast amplification—that is, words used to increase a sentiment such as "very," "deeply," or "extremely"—as well as negation, words used to change the sentiment of the word that follows such as "no," "not," or "never." While natural language processing has inspired a variety of techniques to account for negation and amplification, I opt for simplicity. For words associated with negation, I reverse the sign on the word that follows. This ensures the phrase "not helpful," for example, scores with a negative tone, preventing "helpful" from being counted as a positive word. For words associated with amplification, I add additional emphasis for the word that follows. The phrase "very admirable," for example, scores as two positive words. This approach is similar to that of Godbole, Srinivasaian, and Skiena.

Admittedly, this approach will miss amplification or negation in a more complex format. For example, the phrase "never been admirable" would generally be understood to have a negative tone. In this article's approach, the phrase would be classified as positive, because the word "admirable" is positive and "never" appears to negate the word "been,"

Chart 5
Tone Indexes Using Different Dictionaries

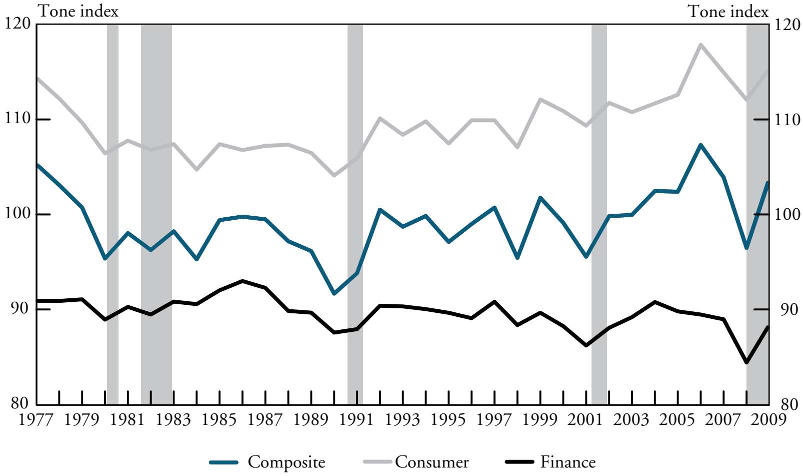

Notes: Index is (positive tone - negative tone) * 100 + 100 where 100 is neutral. Gray bars represent NBER-defined recessions.
Sources: National Bureau of Economic Research, FOMC, and author's calculations.

which has no semantic value. Nevertheless, this approach should accurately capture most instances in which the calculations are affected by amplification or negation.

Both the consumer and finance dictionaries have their strengths as well as their weaknesses, and neither is the obvious choice for this particular investigation. Sadique and others, for example, do not employ Loughran and McDonald's finance dictionary in their investigation of the Beige Book, asserting the text is sufficiently different from 10K filings for the dictionary to be useful.[4] A similar case could be made for the transcripts, as could a similar comparison of the transcripts to online product reviews. To try to achieve some balance in interpretation, I employ both dictionaries for this exercise and use a composite measure of tone that draws equally from them to label the comments. To do this, I score each comment twice and then use the resulting 229,824 labeled comments to calculate a composite tone measure.

The tone measures vary quite a bit across time and dictionaries. Chart 5 displays tone indexes for all three dictionaries. The consumer dictionary measures the tone of the transcripts as consistently positive over the entire period; the finance dictionary, on the other hand, classifies the discussions as consistently negative. The composite measure sits reasonably between the other two measures and shows a similar cyclical pattern, with the tone of the policy discussion seeming to hit a trough just before a measured recession.

Diction and role of speaker

Unlike the FOMC meeting minutes and statements or the Beige Book, the transcripts identify a speaker for each comment. Using the speaker names, I classify each comment as belonging to a Governor, Reserve Bank President, or Federal Reserve staff member. I classify speakers by their respective role at the time of the comment. For example, comments Janet Yellen made from August 1994 to February 1997 contribute to the governor's tally, as she was on the Federal Reserve Board at the time; comments she made from June 2004 through the end of the currently published transcripts in December 2009, however, are counted in the Presidents' comments, as she was then President of the San Francisco Fed.[5]

The relative contribution of different classes of speaker has changed over time. Chart 6 shows the number of comments per speaker that each of the three different classes of participants made each year. Governors contributed the majority of the comments throughout the period, with their contributions peaking in the early 1980s, declining steadily until 2005, then climbing back toward the previous peak. Presidents consistently contributed more comments per speaker than the staff—however, this may not be surprising given the large number of rotating staff members who attend only occasionally.

The method of expression varies across speaker class both in the total number of words and the number of unique words used. Table 1 shows the Governors use a smaller variety of words per comment than either the Presidents or staff. Governors also have the shortest comments, likely due to a larger proportion of questions, which are usually short, instead of longer descriptions of current economic conditions in a district or a prepared presentation on a specific topic.

The measure of tone by speaker class shows a cyclical pattern, with the tone index generally rising during expansions and falling during contractions. Chart 7 shows marked cyclical variations for all

Chart 6

Distribution of Comments across Speaker Class

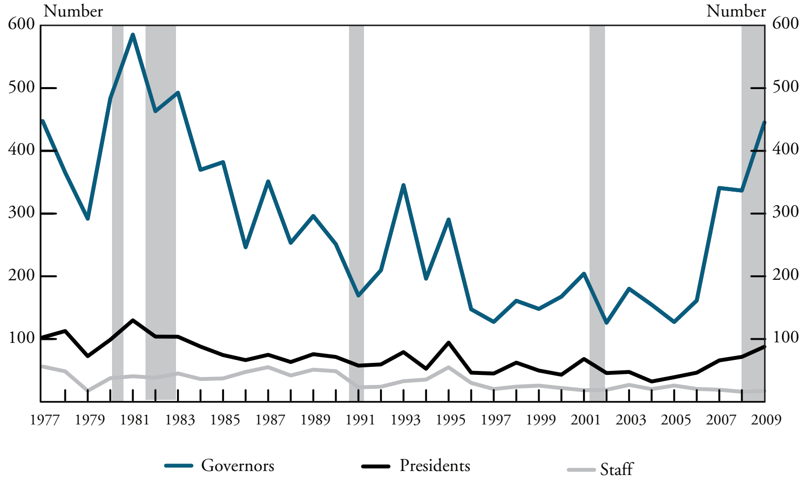

Note: Gray bars represent NBER-defined recessions.
Sources: National Bureau of Economic Research, FOMC, and author's calculations.

Table 1

Word Counts by Speaker Class

	Governors	Presidents	Staff
Average number of words per speaker	47,879	35,467	5,050
Average number of unique words per speaker	582	469	77
Average number of comments per speaker	1,652	601	93
Average number of words per comment	29	59	55
Average number of unique words per comment	0.35	0.78	0.83

Sources: FOMC and author's calculations.

classes but at different levels and different variances. The tone of Bank Presidents, for example, has been consistently more positive than that of the Governors and staff for most of the period. The staff tone has also been consistently more positive, with smaller variation, than the Governors until recent years. Other research has also noted differences in the focus or forecasts among the Governors, Presidents, and staff, and my results seem to align with those findings.[6]

Chart 7

Tone Indexes for Composite Measure across Speaker Class

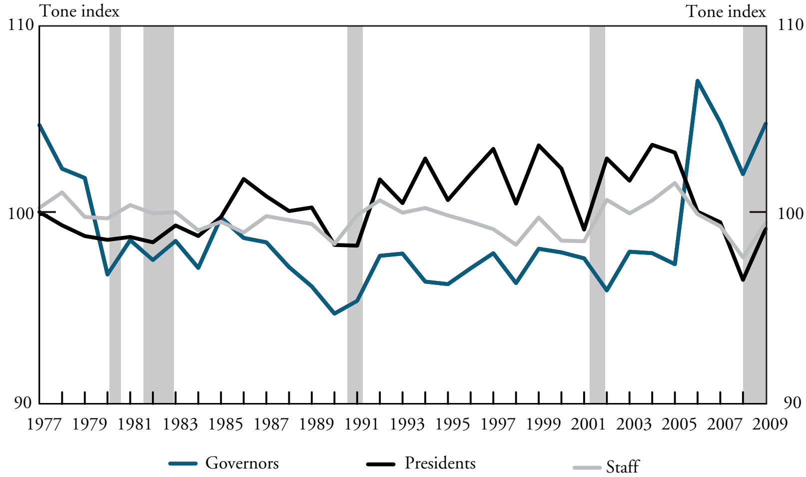

Notes: Index is (positive tone - negative tone) * 100 + 100 where 100 is neutral. Gray bars represent NBER-defined recessions.
Sources: National Bureau of Economic Research, FOMC, and author's calculations.

III. Relationship Between Tone, Transparency, and Real Economic Activity

While the tone measure, both overall and by speaker classification, appears to move in tandem with the business cycle, it may also be related to specific measures of economic activity. To examine whether tone, speaker role, or comment variety are linked with indicators of economic growth or performance, I calculate correlations of the Chicago Fed National Activity Index (CFNAI) with several discussion descriptors and the tone index. The CFNAI is a weighted average of 85 activity indicators constructed to have a mean of 0. It is useful for this comparison because a positive number indicates growth above trend, whereas a negative number indicates growth below trend.

Effect of speaker class

The exercise reveals strong correlations between the tone of the discussions, diction of the participants, and economic activity. Table 2 shows the contemporaneous correlations between several aspects of the transcripts and economic activity as measured by the three-month moving average of the CFNAI. The overall correlation between

Table 2
Contemporaneous Correlations with CFNAI

Variable	Overall	Governors	Presidents	Staff
Total number of words per speaker	-0.56***	-0.64***	-0.43***	-0.05
Total number of unique words per speaker	-0.47***	-0.65***	-0.33*	0.25
Proportion of "inflation" mentions	0.02	0.08	0.0	-0.08
Proportion of "unemployment" mentions	-0.09	0.19	0.01	-0.01
Proportion of "growth" mentions	0.41**	0.38**	0.44***	0.17
Tone measure	0.26***	0.11*	0.28***	0.19***

*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Sources: FOMC and author's calculations.

various features of the Committee discussions and activity is negative and significant: when growth is above trend, the discussions are shorter and contain fewer unique words. Conversely, when economic growth is below trend, the Committee discussions are wordier with more unique expressions. In addition, the relationship between FOMC tone and real economic activity is positive and significant—that is, positive tone in the FOMC discussions today is correlated with a high measure of economic activity.

These relationships hold somewhat when broken down by speaker class as well. The measured correlation for number of words used is lower for Presidents than for Governors, suggesting their contributions to the Committee discussions tend not to decrease as much with a decrease in real activity. Interestingly, the correlation between activity and the expression measures for the staff are not statistically significant. In addition, only one word correlation— "growth"—measures significantly and just for Presidents and Governors. Perhaps not surprisingly, they increase their discussion of economic growth when the economy is experiencing above trend growth. For all speaker classes, the correlation between tone and real activity is positive and highly significant, with the strongest relationship holding for the tone of the presidents.

Of course, correlation does not imply causation, so I cannot concretely determine if the FOMC discussions were positive because real activity was high. While the contemporaneous correlations are strong, both the FOMC discussion tone and the aggregate measures used for comparison may be reacting to the same current market conditions. This does raise the question of whether correlations across time might reveal whether the FOMC discussions lead or lag real activity. Cross-correlation functions show the correlation of the composite tone measure with leads and lags of the economic activity variables up to 12 months.

The correlation between the FOMC's tone and future (or past) economic activity indicates the extent to which the mood of the Committee discussion leads (or lags) real economic growth. Chart 8 shows the cross-correlation of the composite tone measure with the CFNAI for a span of two years: 12 months leading and 12 months lagging. The contemporaneous measure is represented by a lag equal to 0. As the correlation coefficients in Table 2 show, the tone is positively correlated with the CFNAI contemporaneously and in fact leads that index by as much as nine months: for example, a positive tone to the FOMC discussions in January through September is correlated with a positive measure for national activity in October. One interpretation for this long lead time is that FOMC participants have information, forecasts, or expectations that yield a positive tone to the discussion months before the economy experiences above-trend growth. The converse would then also be true: for example, a negative tone to Committee discussions would precede below-trend growth by several months.

As is the case with the tone level and different expression measures, the relationship between tone and activity differs across the speaker classes. The panels in Chart 9 show the cross-correlation of the tone measures with activity for the various speaker classifications. The Governors' tone is positively correlated with economic activity with just a one-month lead. The relationship between the Presidents' tone measure and the activity index is the strongest of all three speaker types and clearly leads the activity measure: a positive tone leads high measured activity by as much as a year. The correlation of the staff tone with activity is positive and significant for longer than the Governors' tone, but does not hold as long as for the Presidents.

The differences in the timing and duration of the effects are interesting in that they vary significantly across the speaker classes. As staff is likely to work more closely with the economic forecasts and other forward projections, it may have relevant information earlier than other speaker types and keep the focus on the periods ahead. Presidents' regional information and strong ties to local business and community

Chart 8

Cross-Correlation of Tone Measure with CFNAI

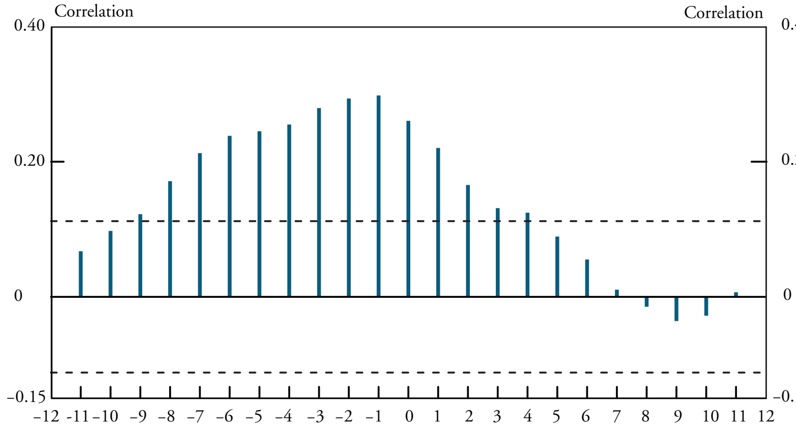

Note: Dashed lines represent 95 percent confidence intervals.
Sources: FOMC and author's calculations.

leaders could also give them earlier information than other classes of speakers, thus contributing to the timing of their tone in the discussion.

Effect of publication

Authors such as Meade and Acosta have posited that a fundamental change in the communication style of the FOMC occurred starting in late 1993. While the meetings were transcribed from recordings beginning in 1976 to compose the minutes, it is not clear FOMC participants were aware of these recordings—or that they expected the transcripts to be made public. In response to a congressional hearing in late 1993, the Federal Reserve Board decided to publish the transcripts from the historical and future recordings with a five-year lag. Meade and Stasavage note that "since 1993 there has been an increased tendency for Committee members to present … pre-prepared statements," which may result in changes in the distribution of words used as well as a change in the general tone measures. Following their work, the analysis here omits the 1993 observations due to possible confusion over who knew about the recordings at what point in the year. Thus, I break the sample into the pre-publication period of 1977–92 and the post-publication period of 1994–2009.

Several measures appear to have changed in the post-publication period. Table 3 shows the differences in the measures of expression

Chart 9

Cross-Correlation of Tone Measure by Speaker Class with CFNAI

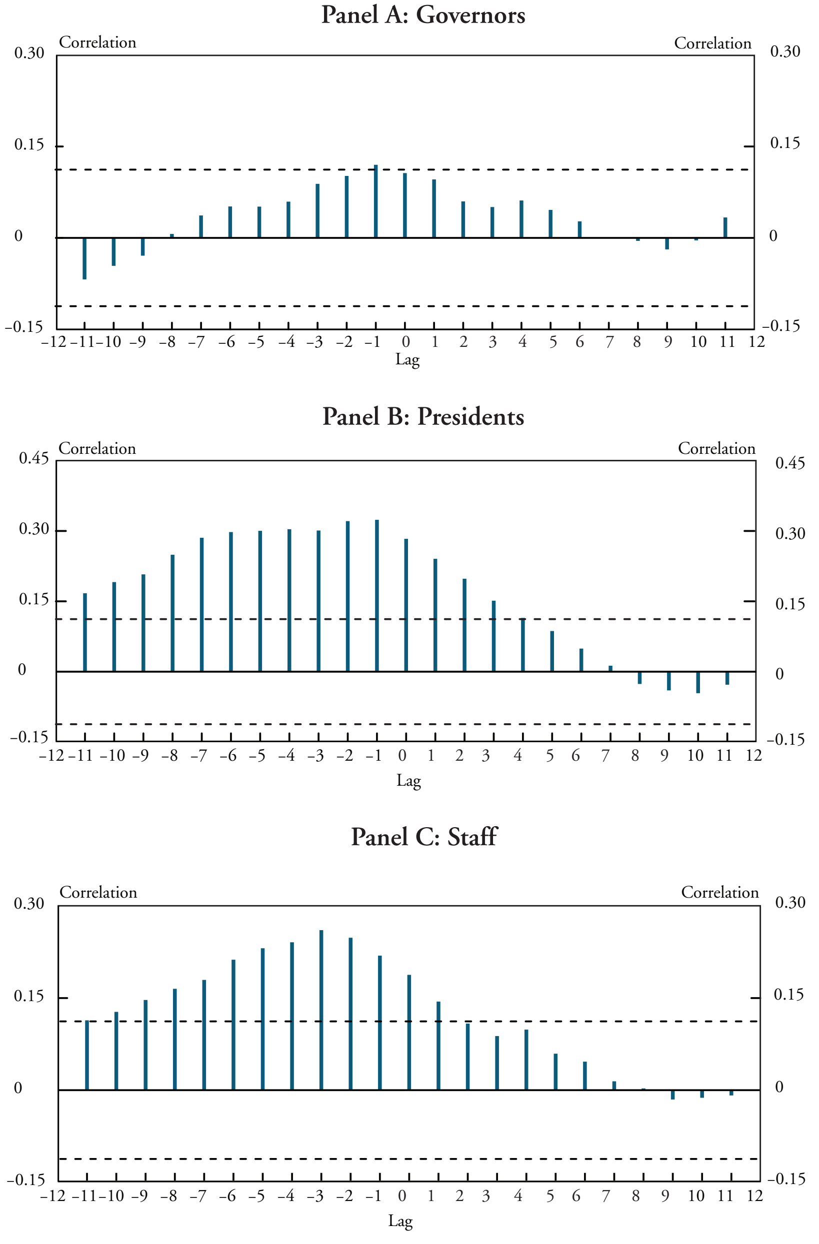

Note: Dashed lines represent 95 percent confidence intervals.
Sources: FOMC and author's calculations.

Table 3
Word Counts by Publication Regime

	Pre-publication	Post-publication
Average number of words per meeting	12,731	13,102
Average number of unique words per meeting	162	115
Average number of comments per meeting	333	302
Average number of words per comment	38	43
Average number of unique words per comment	0.49	0.38

Sources: FOMC and author's calculations.

across the two time periods. The number of comments per meeting decreased, but the number of words used per comment increased—thus, the total number of words was higher post-publication. However, these words were less varied, as both the number of unique words per meeting and per comment declined in the post-publication period. One reason for the differences may be more carefully worded responses or scripted presentations with fewer common words than would be found in less constrained discourse—indeed, former Kansas City Fed President Thomas M. Hoenig said publication "has had some chilling effect on our discussions. I see a lot more people reading their statements. I think it is harder to be as candid as some of us might otherwise be" (Board of Governors 1995).

As before, these word counts should not be interpreted as indicating topic importance but do highlight a marked difference in the diction across the two periods. This result would support Meade and Stasavage's conclusion that discourse did change after the publication of the transcripts became a known and regular occurrence. The change in the form of expression may also support Acosta's finding that speakers had greater conformity—that is, the words they used were more similar—in the post-publication period than in the pre-publication period.

In addition to differences in the number and choice of words across the two periods, the relationship between tone and the CFNAI also differs. Chart 10 shows the relationship between the composite tone measure and the CFNAI across the entire sample period. The correlation appears to be quite close, as the static correlation measure would imply, but the nature of the relationship seems to have shifted over time: the sentiment measure appears to lag economic activity in the early and more recent periods and lead the index in the intervening years.

Chart 10
Tone and Activity Indexes

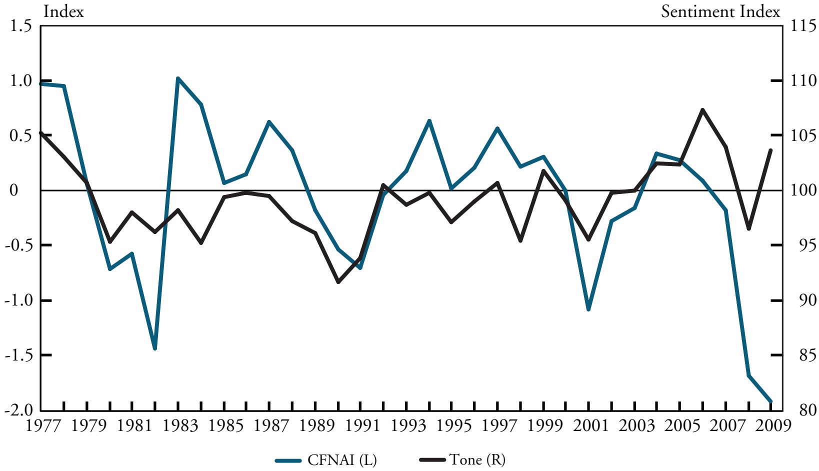

Sources: FOMC and author's calculations.

This change in the relationship from leading to lagging is more visible in the cross-correlation functions. Panels A and B of Chart 11 show that the strength and relationship between the tone of the discussions and real activity after publication lacks the consistent and positive relationship that exists prior to 1993. In fact, any leading relationship for the discussion tone and the CFNAI completely disappears in the post-publication period, and the timing becomes one of a lagging relationship.

A simple regression of the tone measure on CFNAI in each time period shows a measurable change in tone in the later time period.

Pre-publication:

$$\mathit{Tone} = \underset{(0.15)}{99.17} + \underset{(0.27)}{1.02}{*}\mathit{CFNAI}$$

$R^2 = 0.07$

Post-publication:

$$\mathit{Tone} = \underset{(0.18)}{100.11} + \underset{(0.20)}{0.4}{*}\mathit{CFNAI}$$

$R^2 = 0.01$

The intercept implies that at trend growth (CFNAI = 0), the general tone was slightly higher in the latter period, but the correlation between activity and tone was much lower. The coefficient for CFNAI

Chart 11

Cross-Correlation of Tone Measure with CFNAI

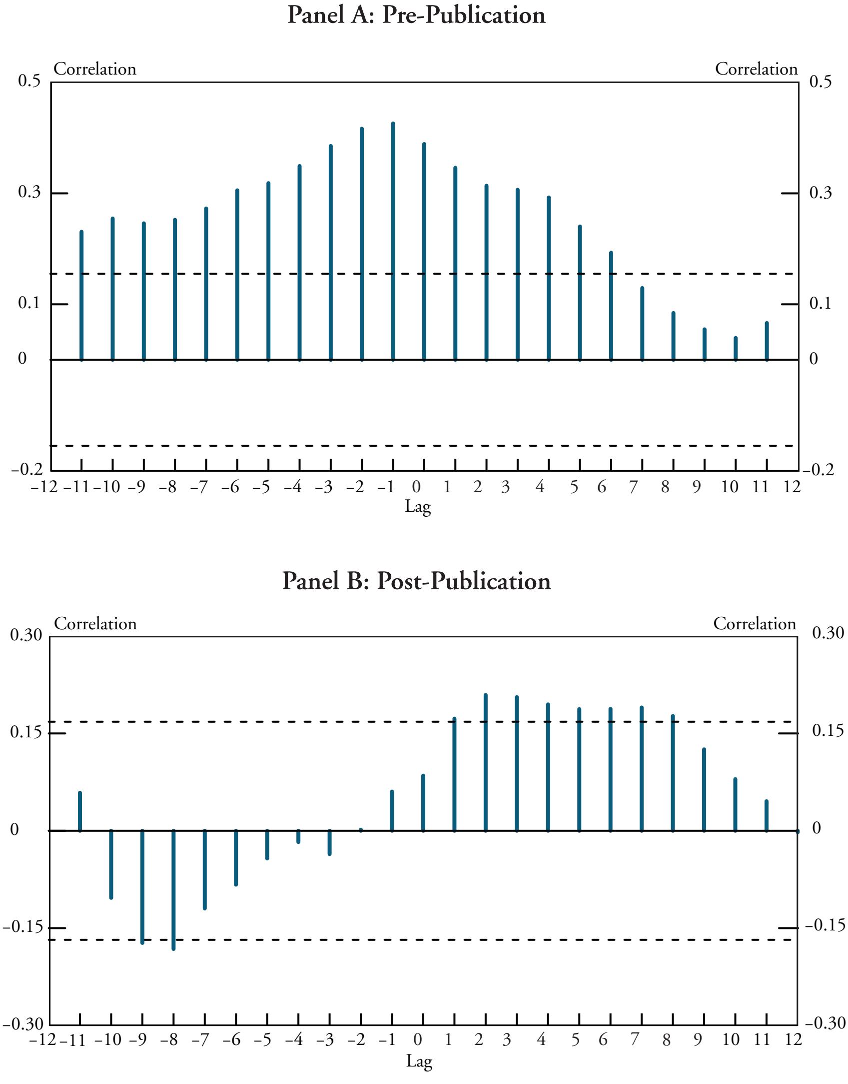

Note: Dashed lines represent 95 percent confidence intervals.
Sources: FOMC and author's calculations.

shows a modulated or dampened effect on the tone of the Committee's discussion relative to the state of the economy after the publication of the transcripts: positive activity sparked a less positive tone in FOMC discussions post-publication than pre-publication.

To explore the differences across speaker class, I separate out the Bank Presidents' and staff members' tone from that of the Governors before and after the publication change.

Pre-publication:

Governor: $\underset{(0.37)}{Tone = 98.23} + \underset{(0.40)}{2.17^{*} CFNAI}$

$R^2 = 0.16$

President: $\underset{(0.20)}{Tone = 99.30} + \underset{(0.21)}{0.72^{*} CFNAI}$

$R^2 = 0.07$

Staff: $\underset{(0.15)}{Tone = 99.97} + \underset{(0.17)}{0.15^{*} CFNAI}$

$R^2 = 0.01$

Post-publication:

Governor: $\underset{(0.36)}{Tone = 98.75} - \underset{(0.39)}{1.26^{*} CFNAI}$

$R^2 = 0.07$

President: $\underset{(0.29)}{Tone = 101.72} + \underset{(0.31)}{1.72^{*} CFNAI}$

$R^2 = 0.19$

Staff: $\underset{(0.2)}{Tone = 99.86} + \underset{(0.21)}{0.75^{*} CFNAI}$

$R^2 = 0.08$

The relationship between tone and activity differs markedly from the pre-publication to post-publication periods across speaker class. Governors exhibit the most drastic change: the relationship between tone and activity switches signs in the post-publication period, indicating that a positive tone correlated with above-trend growth changed to a negative tone correlated with above-trend growth. The Presidents' tone, on the other hand, remained positive in the post-publication period but had a larger coefficient than in the pre-publication period. This suggests Presidents' tended to use a more positive tone when the economy experienced above-trend growth in the post-publication period than they did before the transcripts were published. The staff tone differs across the periods as well: prior to publication, staff tone and activity had no statistically significant relationship, but the relationship became significant in the post-publication period.

IV. Conclusion

The FOMC meeting transcripts provide a unique record over time of monetary policy meetings, yet they have been studied far less intensively than FOMC press releases and meeting minutes. Even with a five-year publication lag, the transcripts are a rich source of detailed information about monetary policy deliberations from which much can be learned. The analysis in this article shows that the tone of the FOMC's discussion varies by speaker class, and that Bank Presidents contribute to the discussion in significantly different ways relative to Governors or Federal Reserve staff members. In addition, basic sentiment analysis shows the tone measure for the Committee discussions is strongly related to real economic activity, but that the relationship varies by speaker class. Finally, the analysis confirms the findings of other research that FOMC discourse shifted measurably after the decision to publish the transcripts in 1993 with both the tone and expression of the discussions changing measurably in the latter period.

These findings suggest that the Committee dynamics and the role of the participants in the meetings are fluid. Much research in this area has examined the transcripts as a single large corpus, without considering the variation over time and across speakers. Adding a time dimension to further text analysis, beyond examining the text before and after the 1993 publication decision, may give even more insight to how policy is formed.

Endnotes

[1]See Danker and Luecke for details of FOMC communications.

[2]See Boukus and Rosenberg; Acosta for examples using a methodology called Latent Semantic Analysis on the FOMC minutes and transcripts.

[3]The publicly available transcripts for one conference call and four meetings are missing pages: July 1979, March 1981, March 1984, November 1984, and August 1992. At the time of analysis, no transcript was available for the November 18, 1980 meeting.

[4]The Beige Book contains regional reports of conditions from the Federal Reserve districts and is an input to FOMC discussions but not an output from the Committee.

[5]I define the "staff" designation as non-Governor, non-President rather than matching comments to actual staff attendees. Therefore, in some instances, especially in the early transcripts, a comment may not be definitively attributable and may appear as "Speaker Y or Z." I classify those occurrences as staff observations unless there is clear indication they should be in one of the other two categories.

[6]For example, see Romer and Romer for a discussion of FOMC versus staff forecasts. For a discussion of Board versus Bank outlooks, see Meade; and Eijffinger, Mahieu, and Raes.

References

Acosta, Miguel. 2015. "FOMC Responses to Calls for Transparency," Board of Governors of the Federal Reserve System, *Finance and Economics Discussion Series* 2015-60, available at http://dx.doi.org/10.17016/FEDS.2015.060.

Board of Governors of the Federal Reserve System. 2014. "Transcripts and Other Historical Materials," *Federal Open Market Committee*, available at www.federalreserve.gov/monetarypolicy/fomc_historical.htm.

———. 1977–2009. "Transcripts and Other Historical Materials," *Federal Open Market Committee*, available at www.federalreserve.gov/monetarypolicy/fomc_historical.htm.

Boukus, Ellyn, and Joshua V. Rosenberg. 2006. "The Information Content of FOMC Minutes," Federal Reserve Bank of New York, available at www.nyfedeconomists.org/research/economists/rosenberg/Boukus_and_Rosenberg_072006.pdf.

Danker, Deborah J., and Matthew M. Luecke. 2005. "Background on FOMC Meeting Minutes," Board of Governors of the Federal Reserve System, *Federal Reserve Bulletin*, vol. 91, no. 2, pp. 175–179.

Eijffinger, Sylvester, Ronald Mahieu, and Louis Raes. 2014."Hawks and Doves at the FOMC," Centre for Economic Policy Research Discussion Paper No. 10442.

Fuksa, Michel, and Didier Sornette. 2012. "The Sentiment of the Fed," *Swiss Finance Institute Research Paper* no. 13–01.

Godbole, Namrata, Manja Srinivasaiah, and Steven Skiena. 2007. "Large-Scale Sentiment Analysis for News and Blogs," *Proceedings of the International Conference on Weblogs and Social Media* (ICWSM). Available at www.icwsm.org/papers/3--Godbole-Srinivasaiah-Skiena.pdf.

Hansen, Stephen, Michael McMahon, and Andrea Prat. 2014. "Transparency and Deliberation within the FOMC: A Computational Linguistics Approach," *Centre for Economic Performance* Discussion Paper No. 1276.

Hu, Minqing, and Bing Liu. 2004. "Mining and Summarizing Customer Reviews," *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (KDD-2004), pp. 168–177.

Loughran, Tim, and Bill McDonald. 2011. "When is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks," *The Journal of Finance*, vol. 66, no. 1, pp. 35–65.

Meade, Ellen E. 2010. "Federal Reserve Transcript Publication and Regional Representation," *Contemporary Economic Policy*, vol. 28, no. 2, pp. 162–170.

———, and David Stasavage. 2008. "Publicity of Debate and the Incentive to Dissent: Evidence from the U.S. Federal Reserve," *The Economic Journal*, vol. 118, no. 528, pp. 695–717.

Romer, Christina D., and David H. Romer. 2008. "The FOMC Versus the Staff: Where Can Monetary Policymakers Add Value?" NBER Working Paper no. 13751, January.

Sadique, Shibley, Francis In, Madhu Veeraraghavan, and Paul Wachtel. 2013. "Soft Information and Economic Activity: Evidence from the Beige Book," *Journal of Macroeconomics*, vol. 37, pp. 81–92. Available at http://dx.doi.org/10.1016/j.jmacro.2013.01.004.